

# The Guide to Business Continuity and Recovery




10100101011011010010101101010110110010101001  
0100100110011001010100011100101010010  
0100101110010010010101001001001  
11101110010101001010101101010101  
10100101011011010010101101010110110  
01001001100110010101000111001010100100001010101

an [internet.com](http://internet.com) Storage eBook

# contents

## The Guide to Business Continuity and Recovery

This content was adapted from Internet.com's Enterprise IT Planet, Enterprise Networking Planet, Enterprise Storage Forum, and InternetNews.com Web sites. Contributors: Richard Adhikari, George Spafford, Kevin Medlin, Drew Robb, Charlie Schluting, and Steven Lewis.



### 2 Disaster Recovery vs. Business Continuity

*Drew Robb*

### 4 Disaster Recovery: Five Things You Might Have Overlooked

*Charlie Schluting*

### 7 Proving Your Disaster Recovery Plan Works

*Steven Lewis*

### 9 Ease Disaster Recovery With SAN Booting

*Charlie Schluting*

### 12 Disaster Recovery and Continuity for the Database Administrator

*Kevin Medlin*

### 18 The Trouble with Virtual Disaster Recovery

*Richard Adhikari*

### 20 Do Your Business Continuity Plans Cover Home Workers?

*George Spafford*

# Disaster Recovery vs. Business Continuity

By Drew Robb

Many IT departments think disaster recovery (DR) and business continuity (BC) are the same thing. As a result, the focus on these subjects tends to rely largely on technology.

And that's a problem, according to Michael Croy, director of business continuity at Forsythe Technology Inc., a Chicago-based IT consultancy and infrastructure firm specializing in BC and risk management.

"Many people are still confused by the terms DR and BC," says Croy. "It is critically important that the DR plan is based on a solid BC plan that has taken into account the reality of the business requirements for recovery. If the DR plan cannot meet the requirements of the business units, it is of no value."

Croy says business continuity plans touch all functions of a business -- from personnel to facilities to IT. In terms of a hierarchical view, business continuity is at the top. Below it is the disaster recovery plan. And under

that come technologies, such as enterprise backup, recovery, and restoration.

But true disaster recovery extends much more broadly than backup processes by using mirrored sites and replicated data to respond to an event. Similarly, business continuity goes well beyond disaster recovery by encompassing every aspect of company operations that could be impacted by a situation. Human resources, power supply maintenance or backup, transportation, food, health, and safety issues all fall within business continuity.

The IT department with its disaster recovery plan is one element of a larger business continuity scenario.

John Glenn, a certified business continuity planner based in Clearwater, Fla., agrees that IT administrators need to take a wider view.

"Most people, especially MIS/IT folks, think BC is just a



The IT department with its disaster recovery plan is one element of a larger business continuity scenario.

new name for DR," says Glenn. "The difference is that DR for IT focuses solely on IT, and what IT perceives as the business unit's requirements. BC, on the other hand, should focus on the business units and, by extension, all the resources required by the business unit."

Industry observers say it's clear that disaster recovery is one element of business continuity. While IT is junior to BC as a whole, the IT organization plays a central role in business continuity.

"It's a big mistake to think the IT department is the only department needed to develop, test, and recover the business," says Gartner analyst Roberta Witty. "It is advisable to form a business continuity program with a dedicated team of people with a senior management sponsor."

IT, though, would provide one representative to the core BC committee.

According to Witty, the committee would be comprised of anywhere from two to five members, depending on the size of the organization. This group would take a wide view of potential disasters.

For example, consider employee health and welfare during an event. In a regional outage, you can't expect personnel to show up for business recovery if they are having serious problems at home related to the event. You must support them and help employees be better prepared at home for disastrous events. The American Red Cross, she says, can be brought in for this kind of training and awareness building.

Michael Gruth, head of system and network support at Deutsche Borse AG, the German exchange for stocks and derivatives, says the IT staff tends to find it easier to relate to the hardware, software, and networking components of DR. He has assembled an Alphaserver/OpenVMS cluster over two sites five kilometers apart. In the process, he discovered there is a lot more to DR than additional Alphas and switches.

"Do not forget things like having an office at your mirror site for remote management," says Gruth. "Also, don't forget the human factor. While it may sound harsh to think about having additional employees to recommence business in the event of a tragedy, this is the reality we live in since 9/11."

To help IT come to terms with a broader scope than disaster recovery, some IT organizations are dropping the term in favor of business continuity.

"We have gotten away from the term 'DR' as it assumes the facility is not available," says Jeff Russell, CIO of The Members Group, an Iowa-based company that provides card processing and mortgage services to credit unions. "BC, on the other hand, deals with how we continue despite business interruption."

### State-of-the-Art Pencils

Disaster recovery projects can easily run aground or fail to be funded if they are done in isolation. Glenn says it is essential to begin every initiative from the business continuity perspective in order to give technology its correct business context.

"Every organization I know about puts BC/DR under the IT umbrella," says Glenn. "My preference is to put BC -- of which DR is a subset -- under the CFO, CEO, COO... someone with some real clout."

To make his point about business continuity not being a matter of technology, Glenn enters the debate about what is the best platform for disaster recovery, or what technological elements are most critical. Should you use OpenVMS or UNIX, mirroring or disk-to-disk backup, SAN or NAS, or all of them? Glenn cuts through the complexity and vendor hype with a simple answer. "My number one DR or BC technology is pencil and paper," he says. "Seriously, it's not about platforms or technologies." ■

# Disaster Recovery: Five Things You Might Have Overlooked

By Charlie Schluting

Disaster recovery plans are useful for more than just audit compliance: If carefully constructed, they can actually work. Purchasing a newfangled Oracle database replication license, for example, is a very small, almost irrelevant part of the overall disaster recovery plan. Let's talk about some of the things disaster planners frequently overlook.

## 1. Remote Is Never Remote Enough

Many datacenter managers in the NYC World Trade Center thought that replicating important functions to the neighboring tower was good enough. If a power failure, water incident, or even a fire threatened one tower, the other could keep operating. Clearly that wasn't remote enough. Perhaps, then, keeping your remote site across town is prudent? No, Hurricane Katrina showed us that was not good enough either.

Think on political and geographical scales. There may be a mountain range that spans two states and creates a scary valley where flooding could occur. More simply,

you may just want to ensure that only one site is located next to an ocean or within an earthquake zone.

Remote sites can be cold, warm, or hot. A hot site is just another datacenter that runs all the time, and is capable of immediate failover to take over duties of a failed datacenter. A full secondary datacenter generally

has a full copy of all data, which means you need to duplicate all of your costs when thinking about SAN or server expansion. Warm sites, on the other hand, often have running equipment, but lack immediate or full failover capabilities. The most common use of a warm site is to ensure the most critical services for a business are kept running. When time and conditions allow, staff can travel to the warm site and start bringing up less critical services

and restoring data from backups. Cold sites may already have equipment in racks, but it's powered off. Staff must travel to the cold site and restore services manually.

Some form of a warm site is generally the best bang



The most common use of a warm site is to ensure the most critical services for a business are kept running.

for the buck, depending on how critical your IT services are. Just remember not to underestimate the severity of a disaster. The next city over is probably too close for a remote site.

### 2. People-Planning: Takes More Time Than Servers

Servers, databases, and storage all have mechanisms to replicate themselves, if you research and choose the right products for the job. It's important to test a data-center, or some critical application's failover procedures and verify that they can run from the remote location. For people whom already have a remote site, the tendency is to stop there, thinking they can survive anything.

Computers don't run themselves, unfortunately.

Your people need to get to the remote site, and if your budget doesn't allow for a fully replicated, millisecond failover, redundant datacenter, your people will be spending a lot of time at the remote site setting things up. Where will they sleep? How will they eat? How will they get there in the first place?

Also, do not assume that critical employees who work in the main datacenter will be able to travel to the remote site and work. Some employees may have family to worry about, some may be unwilling to travel (or work) during an emergency, and some may even be injured themselves. It's important to let each employee deal with personal matters in his or her own way. If an employee goes "missing" following a disaster, but returns after a few weeks, they should be welcomed back with open arms. People deal with emergency situations in their own way, and don't worry, you'll have plenty of staff who are able to help out in a disaster. This is where documentation is critical, since some key personnel may be unavailable and less familiar staff will have to take up the slack.

### 3. Use Your Friends

This isn't really a huge secret, but we're pointing it out because it's often not the first thing disaster planners think about. Shared warm site agreements between businesses are the best way to develop a remote data-center. Often you can work out an exchange, "you can use two racks in my datacenter, and I use two in yours." This doesn't usually work out between competitors, but

channel partners and other business units within your own company are frequently more than happy to exchange space.

Your colleagues in different industries are also an excellent resource. There's rarely any queasiness when you mention storing your data at completely disparate businesses, and often that is the best option. Also, as industries tend to group together geographically, sharing datacenter space with friends in diverse industries will naturally help with "being remote enough."

### 4. You Cannot Imagine Every Scenario

It's very tempting to start thinking about specific disasters, and then draw up plans about how to deal with each one. It's easier that way; if you know a flood has happened, you can conjure up a plan to deal with closed roads and the like. The probability of one of your disaster scenarios being realized is pretty small, unless you spend months thinking about every possibility.

In reality, the best practice for disaster planning is to take a few steps back, and plan for something extremely large. If you have a plan to deal with a completely unreachable city, then you've covered all cases where a subset of the city may be experiencing a disaster. If it happens that people can travel freely, then the plan can be short-circuited and made more efficient. Brief, "if you can" scenarios are far more viable than spelling out what to do in the event of every disaster you can think of.

### 5. Everything You Need is Not There

If you are like most companies, and cannot afford two times (or more) the expense of running your datacenter, you will not have everything you need at your remote warm site. When a disaster recovery plan is activated, scampering employees will often be powering up older machines that may not have been used recently. Hardware may have failed, and even if it doesn't, you still don't have everything you need.

Most companies refresh servers every three years, and it's common practice to extend the life of servers by dedicating phased-out servers as disaster recovery gear. It likely isn't as fast as your current production equipment. Your data has also likely grown since the

last time storage was purchased for the remote site. Perhaps the tape drive your staff was going to use at the remote site to restore data fails part way through. You get the point.

You need a disaster recovery kit, stored at the remote site in a locked container. It should include blank checks, credit cards, phone contact trees, and anything else that a frazzled and hurried employee might need. Again, it all comes back to the people that will execute the plan. ■



# SYMANTEC IS

From antivirus to virtualization. From enterprise data center management to laptop protection. Symantec offers an integrated portfolio of software solutions to help you secure and manage all the assets of your information-driven world. Take control today.

# EVERYWHERE.

[SYMANTEC.COM/EVERYWHERE](http://SYMANTEC.COM/EVERYWHERE)

Confidence in a connected world.



# Proving Your Disaster Recovery Plan Works

By Steven Lewis

Ever since 9/11, we've found an increasing emphasis by top management and government regulators on asking disaster planners to demonstrate that their plans will actually work.

For an organization with even a limited amount of complexity this "show me" requirement can seem overwhelming -- in terms of cost, disruption, and time expended. However, by dividing the task into four levels of increasing difficulty, it is possible to meet that requirement while minimizing disruption.

Basically, there are four kinds of tests available for a contingency or disaster recovery plan: 1) the blink test; 2) audit assessments and structured walk-throughs by "independent experts"; 3) component tests; and 4) "pull-the-plug" exercises.

## Blink Test

This is often linked with the disaster recovery training

cycle. Each task within the plan is first assigned to a specified employee who is then asked to sign a statement saying that they have read and understood the plan as it applies to them and that they are able to carry out their assigned roles within the plan, noting any limitations they may have in carrying them out.



At that point, we've discovered that people begin to speak up and say, "Wait a minute. I'm not authorized for that," or "I don't retain that information," or even "I have family commitments that preclude that." The response has to be, "What do we change in the plan to get you to sign the statement?"

## Audit Assessments / Structured Walk-Throughs

Both internal and external experts next review the plan. Internal experts can include personnel from outside departments who are familiar with how the areas under evaluation operate. These employees are asked

For an organization with even a limited amount of complexity this "show me" requirement can seem overwhelming -- in terms of cost, disruption, and time expended.

to walk through the various scenarios covered by the plan and to provide independent comments, based on their expertise and familiarity with the daily ebb and flow of the specific operations.

To obtain reviews from external experts, representatives of the planning team should be participating in, and to the extent possible, presenting their plan components at region-wide contingency management organizations such as NEDRIX (New England Disaster Recovery Information Exchange), the Business Recovery Managers Association in California, the Business Continuity Planners Association in the Midwest, or others.

### Component Tests

Disaster recovery plans involve many components that can be tested independently. Of course, when the disaster strikes, these components must all work together, but if independent components can be shown to work by themselves, they can be counted on to do their part when the crisis occurs.

Among the specific components that can be independently tested are the recovery and re-installation of backup files; after-hours emergency notification of employees and suppliers; emergency generator operation; and building evacuation procedures.

Component testing also includes simulating a disaster at a single site for organizations that have many locations. By taking one small office offline or relocating it temporarily to its backup site, the department could flush out many problem areas in the transition from normal to crisis-mode and back to normal again. Later on, the plan could be further tested at larger offices.

### "Pull-the-Plug" Exercises

Finally, it is necessary to resolve the question of whether or not all the various plan components can actually work together when they have to. This basically requires a "pull-the-plug" test, in which the entire organization is taken down and then re-opened and operated at alternate sites.

For most organizations, this is simply too disruptive to actually carry out. However, real life often intervenes to make it happen anyway. In those cases, when a mini-disaster happens, planners need to document the events in detail as if it were a test so that afterwards they can assess the following issues:

- Exactly what happened to cause the crisis
- What damages occurred as the crisis unfolded, following the causing event?
- What had been the planned responses to the situation?
- What actions were actually taken by the personnel affected and the responding personnel?
- With the benefit of hindsight, what should have been the responses of the personnel affected and the responding personnel?
- What was learned for the future – what worked and what didn't?
- How should the disaster plan be modified?
- How should the disaster plan modifications be communicated to all personnel?

It is crucial to remember that this testing process is always a work in progress. It needs to be repeated on a regular, ongoing basis, with continual documentation and feedback to all involved. ■

# Ease Disaster Recovery With SAN Booting

By Charlie Schluting

When creating a disaster recovery plan, try not to only think about large-scale disasters. The more realistic disasters, perhaps a failure of an entire blade server, are much more likely to be relevant. In this article, we will talk about the concept of running servers without local disk space, and how SAN and server virtualization can provide extremely flexible recovery solutions.

Enterprises have used disk images for years. Solaris, for example, supports installations via the flash-archive (FLAR) method, whereby a system administrator creates a single image for all similar hardware. The installer is smart enough to copy the archive to local disk, and then modify configuration files to make the server unique.

Without modifications, you cannot simply copy disk images to another server, else a second copy of a server gets brought into existence—with the same IP address. Oh we've all done it—'dd' the disk from one computer to another, then boot it into single user mode and change the host-

name—it's possible, but extremely labor-intensive and tedious.

There is commercial software available to manage this problem seamlessly, especially for Windows server products. They allow an administrator to update a single image, copy it out to all servers, and automatically

update all servers at once. In the Unix world, the story is a bit different, but not all bad news.

## Remote Booting – No More CDs

Unix systems have had remote-boot capabilities for upward of 30 years. Generally a host would request an IP address via BOOTP or DHCP, find a TFTP server (and directory location), and start copying its kernel over the network. Systems that support this type of boot are generally configuring themselves based on network-available information; protocols such as DHCP can provide many

tidbits of configuration information.



The more realistic disasters, perhaps a failure of an entire blade server, are much more likely to be relevant.

To boot a disk image stored remotely, however, a few things are required. First, the image must be unique to each server. Second, the hardware, without any help from a running OS, must support whatever protocols are necessary to obtain and boot from this remote image.

In the first network-boot scenario, the booting servers were dependent on the health of other servers in order to boot themselves. Bootable kernels and the services required to discover their location may be hosted by many different servers, of which a single failure means that other servers failed to boot. SAN-based image booting, on the other hand, requires no other servers. A healthy SAN and properly configured HBA are all that's required.

The idea of storing all your OS disks on the SAN makes good sense for a few other reasons as well. It is extremely efficient for managing many servers at once, and OS images stored on the SAN imply that backups and duplication for disaster recovery can be done via the SAN directly. SAN-based backups (and replication) mean that data is copied off the SAN volumes at the block level, and don't require any backup software to be installed on each server.

SAN booting, as mentioned, requires that each image be customized per-server. This does not mean that each image be custom-crafted. A single golden image may be copied to server-specific storage space for use by that server. When the server boots, it will obviously need to be told something about its IP address, host-name, and other unique attributes. This is a one-time manual configuration that needs to happen, and many enterprises have invented extremely impressive applications to aid in this process. A first-time-boot GUI can prompt the installer for the requisite attributes, but of course, it can also be automated. If the gold image is configured to DHCP and then run some sort of configuration management software, the need to manually configure a new host greatly diminishes.

### SAN Booting, Conceptually

Often people have trouble visualizing how it is possible for a host to do so much before an OS is even loaded. Network booting works because the NIC understands PXE booting, which tells the server where to find and load a kernel. Servers can find their OS disk over a SAN

(or via iSCSI) because the HBA must be configured manually beforehand. Typically, each server is assigned a LUN on the storage that is their boot disk. They can optionally be given other disks too, if required.

The HBA must be configured to know which LUN to attempt to access, and if SAN zoning and the array is configured properly, the HBA can then present a new disk to the server, which can then boot it the normal way. With Sun SPARC hardware you'd set the default boot device to be the device path to the HBA, with additional LUN information. For BIOS-based x86 hardware, there's generally a configuration menu reachable via some unnatural key combination shortly after boot.

### Benefits

Of course, the whole premise of this article is based on disaster recovery. With SAN booting, we can quickly boot a server on new hardware via these few steps:

- Configure new HBA to boot off the desired LUN
- Zone the SAN fabric to allow the new HBA access to the correct storage
- Configure the storage array to allow the new HBA/server to access the LUN

In fact, with SAN virtualization, the last step is not required either. Allocating LUNs and zoning are taken care of by most types of SAN virtualization in a single step.

Using a SAN for DR isn't a new concept; in fact one hardware vendor actively encourages it. HP servers often come with internally accessible USB ports attached to the motherboard. The idea is to put a small thumb drive in each server, pre-loaded with an emergency Linux boot image. If the time comes to reload a server, just boot up, grab a disk image from the SAN, and copy it to local storage. This can be done with either a golden image, or a backed up copy of the actual server's disk. Better, though, is the all-the-time SAN booting scenario. Don't forget, you'll also save the power and cooling costs of two disk drives (at least) per server, which can be surprisingly tremendous.

Quicker than ever, we can recover from server failure, and even boot up disk images in remote locations if we're using SAN-based replication to another array. Perhaps the scariest scenario of all, aside from a true

disaster, is losing a blade server. Some host as many as 16 blades, each with many virtualized guest OS instances running. The death of an entire blade server can literally take out a very large enterprise. Wait, what's that? You're SAN-booting the host OSes on each blade! No problem, just boot them on another server, and you don't even need to touch the hardware or leave your desk. ■

# Disaster Recovery and Continuity for the Database Administrator

By Kevin Medlin

The most important information in most businesses can be found in the database. A lot of time and attention goes into planning for any new database application. Storage, servers, high availability, capacity, and clustering are just some of the considerations.

The same planning process must take place for disaster recovery and business continuity planning of databases. All actions taken to make business critical applications available must be methodical and deliberate. Disruptions are serious events and should not be taken lightly. "It's not about seeing the [recovered] data on your screen, but conducting business," as a 2006 Journal of Financial Planning article put it.

Databases that are at the heart of the business today fall squarely on the critical path of the disaster recovery actions taken when a disruption strikes.

Partial or complete disruptions of a business can be devastating. Business continuity planning can ensure

that capacity is available for critical business operations in the time of need. Practiced professionals in the area of business continuity understand that life and opportunities can continue after a disaster. Understanding the steps involved with keeping a business viable is where some planning is needed.



Destruction of assets can be devastating. Insurance may cover the expense to replace those assets, but it will not put a business back in place overnight. This takes a huge mental and physical toll on workers. These conditions create burdens and stress on employees and their customers. Without a disaster recovery plan in place, there is little hope of ever getting a

business back on its feet.

## Requirements

One of the first things needed are the requirements for each database supported. Recovery times are probably

Databases that are at the heart of the business today fall squarely on the critical path of the disaster recovery actions taken when a disruption strikes.

the most important of these requirements. The difference between a few seconds of downtime and a few minutes of downtime can be quite substantial. Some business units may have a tolerance for a few hours. This must be known for each database for your plan to be effective. "...[Y]ou have to prioritize what you need in order to function... you have to figure out what is actually mission-critical," wrote Charlene O'Hanlon in a 2007 T H E Journal article.

Another important answer needed is in reference to data loss. If little to no data loss is acceptable, then a disaster recovery solution can become a budgetary concern. If the backup from last night will suffice, then this can lead to major cost savings.

Capacity can be a concern at the disaster recovery site. Customers should be asked about performance degradation and what is acceptable. This can be a tricky question to answer, and customers will usually need assistance to figure it out. If left to themselves, they will almost always answer that no degradation is acceptable.

Another question that should accompany performance degradation is finding out about the number of users that will be accessing the system during the disruption. These two answers will help to identify a more accurate capacity. What should be explained is that during the disruption, the entire corporate population may not need access to the enterprise application. Possibly only power users may need the system to run business critical functions for the enterprise.

Human Resources applications are one example. An HR application may be available to the corporate population during normal operations for viewing pay stubs, updating W-2s, and so on. During a disruptive event, these rights could be suspended but power users could continue to run payrolls, enter benefits, hire and fire employees, and the like. It is possible that far less capacity is needed than originally thought necessary, which can mean more databases on the same servers, as long as the databases will not interfere with one another's processing. Virtual servers can be used as well.

"... [Y]ou would re-instantiate the virtual machines at a higher ratio (density) of virtual-to-physical. Consequently, organizations that can tolerate a slight

## Is Your Recovery Plan Good Enough to Save You?

By George Spafford

Unfortunately, not all organizations realize the critical need to internalize disaster recovery planning and may figure they will let the government help them if the time comes. What they don't realize is that even if a disaster strikes, there may not be aid. They must take care to preserve their own business continuity.

Organizations simply must take control of their own recovery plans.

Hurricanes like Katrina and Rita are vivid in peoples' minds, as is the outcry for assistance from the government and private organizations. However, assistance isn't always forthcoming.

In September 2005, Wisconsin was struck by 27 tornadoes that damaged 400 homes. Their request to be declared a federal disaster area to get government assistance was denied.

Can you gamble on getting assistance?

Despite living in a city that was below sea level, many in New Orleans did not have flood insurance, yet were covered for hurricanes – or so they thought. Heated debate and lawsuits arose from carriers declining claims based on arguments that the property damage was not caused by the hurricane directly, which would be covered. Some claim the storm surge and subsequent flooding is what caused the damage and that would not be covered by insurance policies.

The issue is that flooding requires a separate rider that many did not buy. If those families and businesses do not get reimbursed from insurance, how will they fair? Have you checked your insurance policies lately against your most likely risks to make sure you have the appropriate coverage to ensure that recovery is possible?

To worsen many already dire situations, some organizations in New Orleans dutifully sent their backup media to offsite storage sites located around the city. Not only did some groups lose their on-site data, but the offsite data was destroyed, as well.

*continued*

drop in performance can build a much cheaper secondary data center to handle temporary disruptions," according to a Nemertes Research report by Andreas Antonopoulos.

Accessing the databases and applications is another important matter. If the primary place of employment is no longer habitable, employees will need a place to go for office space and workstations. Workstations will need to be equipped with necessary software for database connections. This important point must not be overlooked.

Testing is very important. Determine the frequency with which you will need to test your disaster recovery plans. Only through testing of the plan can issues and problems be discovered and corrected. Testing can also bring opportunities to make improvements to the disaster recovery plan.

Since nothing stays the same in business very long, you will find the same quality in disaster recovery plans. To keep them relevant and up-to-date, testing must become a regular occurrence. Testing may occur yearly, twice per year, or quarterly. The more practical experience individuals can get with the disaster recovery plan and the disaster recovery site, the better off everyone will be during a crisis situation. Familiarity will build confidence in individuals and the equipment and systems they are working on.

Usually, disaster recovery setup is not an emergency. The emergency only comes during execution of the plan. Still, a timeline should be put in place when planning disaster recovery for databases. It is unfortunate that many times, other projects push disaster recovery to the back burner. Make disaster recovery part of all projects so that it can be completed in a timely manner.

Moving back to the primary site will be a joyful time. It can also be quite hectic, since it needs to be done quickly. No one wants to stay at the disaster recovery site any longer than they have to. Plan the return much as would be done with the go-live of a new application. Plan the downtime, migrations, testing, go/no-go decision and fallback procedures. Everything should be scheduled and users made fully aware of the outages and changeover schedules.

There should be someone, or some people, in the

Given your most likely risks, do you have a backup process that safeguards your data from regional incidents? Do you need to guard against regional disasters, and if so, how far away must the backups travel?

### **The Need for Planning**

With just these few examples in mind, when was the last time you and your team sat down and ran through the most likely scenarios that threaten your organization? The careful review should move beyond abstracted risks and focus on layered situations. Move past "what if we lose power?" and instead focus on realistic matters such as "what if lightning takes out both the primary and secondary grids that feed our facility?"

The power company's communication structure is in disarray and an estimated time to recover is not even available. What must be done immediately? What do we do 30 minutes into the outage? What do we do an hour in? At what time do we begin powering down systems and in what order? How do we inform employees?

The idea is to use realistic situations to foster dialogue and to capture and formalize ideas that are scattered through the team. The end result must be a disaster recovery plan that covers the most likely scenarios. Whether there are three, five, or 20 scenarios, the exact count will depend on the organization and the risks that confront it.

The goal is to plan to the level that management feels is adequate.

Whenever a disaster strikes, even a small one, take the time to review lessons learned. Determine what worked well, what did not and revise plans accordingly.

### **Business Continuity**

Moving beyond disaster recovery is the idea of business continuity.

How will you keep the business running during some kind of disaster? If disaster recovery is concerned about restoring a given service back into production, business continuity planning is concerned with the holistic issues surrounding keeping the business running or getting back up and running as quickly as possible to minimize impacts.

Some organizations get hit by a disaster and disappear. We, of course, don't want that to happen to us.

*continued*

organization who will make the decision that a disaster has struck and failover should now take place. Determine who that person is and how the information will be communicated. Ideally, the information will be distributed in multiple forms. Rarely in a disaster will all the normal lines of communication be available to an organization.

There are many key roles that are critical to the success of the database administrator. A server administrator will have to install and set up the server. A system administrator will be needed to install and set up the operating system. A storage administrator will be necessary to duplicate the disks accordingly. Application developers will need to assist with troubleshooting errors detected by the user community. These are some of the people that a database administrator will rely on.

Many, if not all, of these steps can be accomplished prior to any disaster and tested. There can also be problems at the time of failover where some of these areas may need to be revisited. The database administrator may know who to call and work with during normal times, but what happens when a disaster strikes and some primary support personnel are not available? They could be taking care of injured family members or injured themselves. What if your database administrator is not available? Contingencies for these scenarios should be put in place.

It is imperative for employees to know whom to call when they have an issue.

One of the best ways to avoid a situation with availability is cross-training employees. An employee that knows more than one job function can become essential and can play a key role during a disruption by knowing more than one area or job function.

Some people may not be able to make it to the recovery site, leaving some areas not covered, noted Eric Maiwald and William Sieglein in "Security Planning & Disaster Recovery" (McGraw-Hill). The cross-training should not be a complete shift from their normal profession, unless requested by the employee. What is usually better is to have an employee learn a skill that is new, but in the same profession they are currently engaged.

If we return to our power example from above, think about what business processes are most critical to our ability to stay operating. What is needed to operate? If the automated systems are down, can they run manually?

These questions are aimed at understanding the organization's requirements and then layering IT's capabilities in to support the business.

Organizations must review their risks and then develop options to mitigate continuity risks.

For details, there are many resources on the Web that have been quietly evolving. There is a wealth of recommended practices out there to aid in your planning, including recommendations in ITIL and ISO 17799. Furthermore, discuss matters with your team and industry association to get started.

There are many avenues to consider. Groups that haven't dusted off their disaster recovery and business continuity plans since Y2K should get them out and run through them, thinking about the disasters most likely to strike. The scenarios should be detailed enough that responses are gauged, corrective actions defined and investments approved.

Organizations can't take their responses for granted. If they do, they might be faced with the day when planning would have made the difference between being in or out of business. ■

For example, Oracle database administrators can cross-train as SQL Server database administrators. They are already familiar with the concepts, SQL, structures, and other features of database administration. It should mostly be a matter of learning the different toolsets for the new database software.

This can be a win-win for the employee and the organization. The employee learns a valuable new skill that can enhance their career. The organization gains an employee that has multiple skill sets that can be called upon in times of normalcy and times of crisis.

### Backups

Requirements for a database will drive the type of backups you make for it. If a database can have several hours of downtime and the last night backup will work sufficiently, then a full backup will be fine. If little to no

downtime or little to no data loss is acceptable, then full backups will not do the job.

Technologies such as remote mirroring will have to be investigated. In remote mirroring, all changes made to the production system are copied to the disaster recovery site. This is normally considered in an asynchronous context, since most disaster recovery sites are at some distance away from the primary site. When a fail over is called for, databases can be recovered with the mirrored data for business continuance.

Data replication is another technology that can keep disaster recovery databases updated. The native settings of the software replicate changes as they occur from production databases to databases at the disaster recovery site. This can be altered so that changes are applied on a schedule, like every four hours. This would be for a data recovery scenario in case a user made an error. The database administrator could use the data from the disaster recovery database to correct the error in production because the changes had been delayed.

### Installations

Installation of database software should be a fairly routine task for a database administrator. It should also be the same across servers with the same database versions. Installation and setup should be well documented. There is always the possibility that a database administrator will not be available when a fail over is called for. Clear and concise, step by step directions will give technical professionals from another area the ability to stand in for a missing database administrator and set up the database software.

This being said, each production server is different. Certain things may need to be done to prepare the database. Special scripts will sometimes need to run, or jobs to load or unload data. These steps for individual databases and the order in which they should execute also need to be well documented.

### Making Good Use of the DR Site

The best way to set up disaster recovery is by having a dedicated site with servers available and application software running so that an immediate fail over can be done when called for. This approach is also very expensive and not always popular. There are ways to imple-

ment disaster recovery sites, save money and be practical, all at the same time.

An excellent approach for the dual use of just such a facility is testing of upgrades. All operating systems, applications, and databases require regular maintenance patches, fixes and upgrades. With environments available as exact duplicates of production systems, these are prime locations to test the maintenance releases.

Patches and fixes can be applied to a disaster recovery system on a regular schedule. An approved test plan can be administered against the environment to check for issues with the maintenance release. If no issues are found, the patches can be left in place and migrated to the test environment on a regular schedule as well. If no problems are found, the patches can then be migrated into production on a regular schedule.

If any issues are found at the disaster recovery site or in the test system, then the patch can be rolled back or tickets can be opened with the vendors if problems are minor. This eliminates the need for a separate laboratory environment, which can also be very costly. No additional hardware, software, licenses, maintenance, administration, or space would be needed for a lab to test maintenance releases.

If you do not currently have a lab for testing patches and fixes for software, then this can be of a substantial benefit in three areas. The money has already been spent on the disaster recovery site, which was a necessity in itself. Secondly, a duplicate environment of your production systems now exists to test software patching, negating the need for a laboratory. Thirdly, less administrative maintenance is spent on systems once they are patched. Keeping software patched and fixed to current levels reduces downtime and the amount of time administrators spend on system repairs.

This approach can be especially helpful for database administrators. Many times a server may be available for database installations, patching, and upgrades, but rarely are there complete environments for these tasks. The need for application developers and users is to test the application against the database after the patches have been installed. The database administrator can perform some limited testing, but the true tests come

when users put the system through the motions.

Stocking the disaster recovery site with test servers is another great way to get the disaster recovery site up and running quickly and maximize the value of those servers. In most, if not every case, these servers are purchased for every new project that will be migrated into production. Test servers should be purchased with the same specifications, or better, than production. Most test servers will need higher capacity because more databases, application servers, Web servers, and the like will be running on them than the production hardware. With test servers in the disaster recovery facility, much of the work of software installation is already done. Disaster recovery instances can be created on test servers and left idle. Application servers, Web servers, and databases just wait for the day that a fail over will be alerted.

Using virtualized servers can assist in lower costs for a disaster recovery site, particularly as the technology becomes less expensive and less complex. It is now much easier to implement virtual servers than it has been in the past. Today, many applications, operating systems and databases support server virtualization software. This has changed since many of the virtualization vendors have tried to work closely and cooperate fully with the other software vendors.

Pressures from customers have also driven software companies to work with virtualization companies to certify and support their products. Through virtualization, a physical server can be imaged and reproduced in a virtual environment. A production system consisting of a Web server, application server, and a database server can all be imaged and virtualized on a single physical server. This effectively consolidates three physical servers down to one without losing any functionality. Capacity may not be equal, but it may suffice perfectly

in a disaster recovery scenario. This does not mean that all applications will work together on virtual servers; they must be able to coexist.

### Mentoring

A step beyond cross-training is mentoring. A mentoring program allows subject matter experts to work directly with management-identified employees who are interested in becoming experts in a different field than the one they are currently in. This can become a large financial gain for employers while increasing employee morale as well. Mentoring can also work well for employees who wish to cross train to qualify for positions on other technology teams that have unfilled vacancies.

By identifying and opening career opportunities across teams, individuals feel a sense of empowerment and are not stuck in their current roles. For example, a database administrator position may be difficult to fill externally. A current developer with talent, ability and desire to become a database administrator could miss an opportunity to make a lateral move due to lack of experience. Through mentoring, the developer could continue in her current role while cross training in a potentially new career path. In this way, mentoring programs can help manage expected retirements and workflow fluctuations while providing alternative career paths for qualified candidates.

A mentoring program spreads knowledge within and across teams, providing support when subject matter experts are inaccessible or incapacitated, a critical consideration for disaster recovery. By documenting processes and procedures through a mentoring program, the ability to respond quickly to outages or disasters is dramatically enhanced. ■



# SYMANTEC IS

The industry leader in backup, clustering, and replication software across multiple platforms.

# DISASTER RECOVERY.

[SYMANTEC.COM/EVERYWHERE](http://SYMANTEC.COM/EVERYWHERE)

Confidence in a connected world.



symantec™

# The Trouble with Virtual Disaster Recovery

By Richard Adhikari

As enterprises virtualize their data centers to cut costs and consolidate their servers, they may be setting themselves up for big trouble.

According to the latest disaster recovery research report from Symantec, based on surveys of 1,000 IT managers in large organizations worldwide, 35 percent of an organization's virtual servers are not included in its disaster recovery (DR) plans.

Worse yet, not all virtual servers included in an organization's DR plan will be backed up. Only 37 percent of respondents to the survey said they back up more than 90 percent of their virtual systems.

When companies virtualize, they need to overhaul their backup and DR plans, Symantec says; the survey found that 64 percent of organizations are doing so.

"That's no surprise, because virtualization has had a huge impact on the way enterprises do disaster recovery," said Symantec senior product marketing manager for high availability and disaster recovery Dan Lamorena.

So why are virtual servers being left out of DR plans, or, if they're included, aren't being backed up? That's because enterprise IT just does not have the right tools to back up virtual servers, according to Symantec.

The biggest problem for 44 percent of North American respondents was the plethora of different tools for physical and virtual environments. There are so many that IT doesn't know what to use and when.

Another 41 percent complained about the lack of automated recovery tools. Much of the disaster recovery process is manual, although VMware recently unveiled a tool to automate the run book.

Another 39 percent of respondents said the backup tools available are inadequate.



When companies virtualize, they need to overhaul their backup and DR plans, Symantec says; the survey found that 64 percent of organizations are doing so.

Hewlett-Packard, IBM, CA, and smaller vendors such as ManageIQ, Avocent, and Apani offer tools to manage both the virtual and physical environments. And companies like Hyperic are bringing out new tools.

However, virtual server management tools, being relatively new, are not as sophisticated as their counterparts for the physical environment. Also, they have not been around long enough for users to be familiar with them. For example, provisioning, or setting up, virtual machines from physical ones and vice versa can also be a problem, and tools for this have only recently emerged.

"Virtualization makes some aspects of backup and disaster recovery more difficult," said Symantec senior product marketing manager for NetBackup Eric Schou. "IT shops are still struggling with the steep learning curve."

Porting over solutions from the physical environment won't work, Schou said. "IT shops need to get solutions that are finely tuned for virtualization," he added.

### Failing DR

Judging from the results of the survey, IT is still not as familiar with DR as it should be. DR testing is a mess.

A whopping 30 percent of respondents said their DR

tests failed. That's better than the 50 percent failure rate in 2007, but it's still pretty scary.

For 35 percent of the respondents, the tests failed because "people didn't do what they were supposed to do," Lamorena said. This means that much of recovery is still a manual process, and companies must begin looking at automation, he said.

Another cause is that tests are not run frequently enough. That's because "when you run a test, it disrupts employees and customers," Lamorena said. He added that 20 percent of the respondents said their revenue is hurt by DR tests, so "the tests cause the same pain to their customers as if they had a real disaster."

Finally, the survey found that top-level executive involvement in DR planning has fallen. "Last year, the C-level involvement on disaster recovery committees was 55 percent; this year, it's 33 percent," Lamorena said. C-level executives are CIOs, CTOs and CEOs.

Lamorena finds the reduction in top-level involvement disturbing because it could lead to more problems with DR. "That's a huge drop, and we've been thinking about this day and night," he said. "What's alarming is, companies may be getting a little lax and don't think they'll be affected by a disaster." ■

# Do Your Business Continuity Plans Cover Home Workers?

By George Spafford

Organizations concerned with surviving a disaster have undertaken Business Continuity Planning (BCP). Studies have shown that without adequate planning, some organizations will not be able to continue operations. Indeed, BCP should be performed by all organizations with the time and resources invested commensurate with the risks and constraints of the organization.

In cases like the risk of a flu pandemic or something else that would keep workers from reporting to the office, the response of many organizations is to have their personnel work from home utilizing personal, or even company-funded, high-speed Internet connections. While an intuitive approach, there are additional risks to consider as a recent experience helped highlight.

The first risk to consider is that the phone companies, hitherto referred to as "telcos," have a policy of taking care of business customers first. This recently hit me very hard when AT&T shut my DSL down for almost

two weeks due to a simple billing change request coupled with an error in AT&T's nearby central office.

During the course of numerous phone calls, it came to light that even though my firm is a legal corporation, their customer service representative had entered my number as a residence and not a business. Even though I could show that my articles of incorporation clearly identified my home office phone number, it did not matter. I had zero recourse with AT&T and the impact of losing DSL to my business has been bad.

My savior was using my dial-up modem – a technology I haven't relied on for communications for several

years, even during international travel. As a result of my experience, my first recommendation would be to ensure that you review your organization's projected needs during a crisis and then review what your home workers' needs with your telco.



... my first recommendation would be to ensure that you review your organization's projected needs during a crisis...

My message to the telcos is to recognize that more and more people are operating from home and residential service offerings do not suitably recognize the needs of people working from home. For example, consider offering a tiered service plan to better meet the needs of people conducting business from home.

### Have a Failsafe Connection Ready

Second, the episode highlights that people working from home may need more than one avenue to connect to the Internet. A DSL or cable connection may be the first line and then dial-up as a failsafe. There are a number of wireless options now as well, ranging from broadband cards from cell phone companies to wireless Internet Service Providers.

An organization's risks associated with loss of contact should drive expenditures. For example, some users may need two or even three alternatives, including a more direct connection than the public Internet, that are constantly active while others have the hardware and software and have a negotiated activation time. Again, let the risks drive the expenditures.

Third, we are all assuming that in the event of a disaster that the telcos will be operational. In fact, some may outright fail, others have degraded service or slowly come to a halt.

For example, if a pandemic were to hit, forcing people to work from home, how long could telcos, power

companies and other infrastructure providers sustain operations? What are their plans? In preparing for Y2K, we asked our suppliers to share their plans -- often necessitating a Non-Disclosure Agreement being signed if one did not already exist. The point is that we recognized the need that the survival of our firm depended on others surviving as well.

Fourth, during my DSL loss, not only did I lose Internet access, I was even further impacted by sitting on the phone with customer service over and over. It causes one to stop and ponder that in the event of a disaster, how will support be coordinated? Given your location, resources and needs, what can be done to optimize support?

Being without DSL for almost two weeks was equally frustrating and illuminating (one of the side effects of being a prolific writer, I suppose) and given me an opportunity to think about just how important high-speed Internet has become and the potential impacts of telco prioritization rules on corporate business continuity plans.

Organizations need to ensure they have adequate continuity plans in place, that their infrastructure vendors can support the plans, and that the vendors have their own plans for that matter.

Undoubtedly it is better to do the work now and have an idea of what to expect than to have a disaster and enter not knowing the questions or the answers. ■